

PATENT APPLICATION

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re application of

Docket No: Q66255

Robert HOFNER, et al.

Appln. No.: 09/989,377

Group Art Unit: 2667

Confirmation No.: 1293

Examiner: Christopher P. GREY

Filed: November 21, 2001

For: APPARATUS AND METHOD FOR LOAD BALANCING IN SYSTEMS HAVING
REDUNDANCY

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

MAIL STOP APPEAL BRIEF - PATENTS

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

In accordance with the provisions of 37 C.F.R. § 41.37 Appellants are submitting an Appeal Brief to appeal from the Final Office Action dated December 14, 2005 (hereinafter "the Final Office Action"), wherein rejection of claims 1-38 have been maintained. This Appeal Brief is accompanied by a Submission which includes the required appeal fee set forth in 37 C.F.R. § 41.20(b)(2). Appellants' Notice of Appeal was filed on June 14, 2006.

Table of Contents

I. REAL PARTY IN INTEREST.....	3
II. RELATED APPEALS AND INTERFERENCES	4
III. STATUS OF CLAIMS.....	5
IV. STATUS OF AMENDMENTS.....	6
V. SUMMARY OF THE CLAIMED SUBJECT MATTER.....	7
VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL	15
VII. ARGUMENT.....	16
VIII.CONCLUSORY REMARKS.....	<u>24</u>
CLAIMS APPENDIX:	<u>25</u>
EVIDENCE APPENDIX	<u>40</u>
RELATED PROCEEDINGS APPENDIX.....	41

I. REAL PARTY IN INTEREST

The real party in interest is EXANET, INC. (Assignee) by virtue of an assignment executed by the inventors Mr. Amnon A. Strasser and Robert Hofner on November 18, 2001 and filed on November 21, 2001 along with a form PTO-595.

II. RELATED APPEALS AND INTERFERENCES

Upon information and belief, there are no other prior or pending appeals, interferences, or judicial proceedings known to Appellants, Appellants' representatives or the Assignee that may be related to, be directly affected by, or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Each of the pending claims 1-38 are rejected (see Final Office Action dated December 14, 2005).

IV. STATUS OF AMENDMENTS

The Advisory Action dated April 12, 2006, indicates that the Amendment filed March 14, 2006 has been entered. Therefore there are no pending unentered amendments.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

Many modern computer systems require support for load balancing as well as redundancy. Supporting both features as separate entities is costly, especially since redundant units in the system are used infrequently. The present invention seeks to have a system that integrates load-balancing capability as well as redundancy capability.

The present invention relates to computer systems having redundancy capabilities and to techniques for load balancing between units of the redundancy mechanism. According to the invention redundant units within a networked storage system are used for load balancing. Load balancing refers to fine tuning of a computer system to more evenly distribute the data and/or processing tasks across available resources. In a specific example of a clustered system that handles financial transactions, load balancing distribute the incoming transactions evenly to all servers that comprise the cluster.

However, in computer systems that support redundant resources, such resources wait in a standby mode. The redundant resources are activated if an active system becomes inoperative.

In the exemplary implementation shown in FIG. 1, the system 110 comprises multiple terminals 110-1, 110-2, 110-n that are connected through a connectivity medium 120. A plurality of resources 130-1, 130- is connected to the connectivity medium 120. To properly run the system, two separate processes are implemented; namely, load balancing and failure detection followed by correction. During normal system operation, all the resources 130, as well as all the paths in the connectivity medium 120, equally share the probability of being used by system 100. Thus, during normal operation, a nearly even load distribution between all the

resources 130 as well as the paths in the connectivity medium 120 leading to the resources 130 is maintained.

Upon detection of a failure of a path in the connectivity medium 120 by a monitoring system, a redistribution of the load takes place. The failed path is “removed” from the connectivity medium 120, for load balancing. Since at least one or more redundant paths are available in the connectivity medium 120, available overall system “on” performance is maintained. Similarly, when a resource is “down” or otherwise inoperative, tasks that were assigned to the inoperative resource are redistributed to other active resources having properties similar to those possessed by the inoperative resource.

As shown in the exemplary flowchart of Fig. 2A, at S210, the system receives and recognizes a request for accessing a system resource. At S215, a determination is made as to which resource of those resources potentially available for the specific resource request will be made available for use. Once the specific resource to be used is determined, at S220, a determination is made as to which of the paths available in the connectivity medium ought to be used. Once a path in the connectivity medium has been determined, at S225, the system informs the requestor of the selected resource and path in the connectivity medium.

As shown in FIG. 2B, in determining which resources to assign to the resource request, at S250, the least loaded resources of the type required to fulfill the resource request are searched for and designated. At S255, the number of such available resources is determined. If there are two or more resources available for use then, at S260, a selection function is executed in order to

determine the availability of a single resource. In S265, the resource requestor is informed of the specific resource to be used. A person skilled in the art could easily implement both of these methods in hardware, software or combination thereof.

Specifically, the present invention as recited in claim 1 is a system comprising at least one terminal node and at least one network resource. (Specification p. 9, ¶ [29], ll.1-3, Fig. 1, item 130-1) The network resource has at least one redundant matching resource. (Specification p. 12, ¶ [32] ll. 8-18). The system includes a computer that transfers tasks from execution on the network resource for execution to the redundant matching resource if said network resource fails. (Specification p. 12, ¶ [32], ll. 8-18). The computer balances execution loads between the tasks executed by the network resource and the tasks executed by the redundant matching resource. (Specification p. 12, ¶ [32], ll. 8-12). A communication medium connects the computer, the terminal node, the network resource and the redundant matching resource. (Specification p. 11, ¶ [30], ll. 13-18, Fig. 1, item 120). The communication medium has at least two independent communication paths between the terminal node, the network resource and the redundant matching resource. (Specification p. 11, ¶ [30], ll. 13-18).

Independent claim 9 recites a system comprising a plurality of terminal nodes (Specification p.11, ¶ [29], ll. 1-3) and a plurality of network resources (Specification p. 11, ¶ [29], ll. 1-3). A plurality of redundant resources are provided. (Specification p. 4, ¶ [12], ll. 1-4). Each of said plurality of network resources closely matching at least one of said plurality of redundant resources. (Specification p. 4, ¶ [12], ll. 1-4). A computer moves tasks from a failed network resource to a redundant resource that closely matches the failed network resources.

(Specification p. 4, ¶ [12], ll. 4-7). The computer balances execution loads between the tasks executed by said plurality of network resources and the tasks executed by said plurality of redundant resources. (Specification p. 4, ¶ [12], ll. 4-7). A communication medium connects the computer, the terminal nodes, the network resources and the redundant resources. (Specification p.4, ¶ [12], ll. 8-11). The communication medium has at least two independent communication paths between said terminal nodes, said network resources and said redundant resources. (Specification p.4, ¶ [12], ll. 8-11).

Independent claim 16 recites a method for balancing loads of tasks executed in a network system. (Specification p. 5, ¶ [13], ll. 1-2). The network system contains a plurality of communication paths, a plurality of redundant communication paths independent of said plurality of communication paths, a plurality of network resources of differing types each of said network resources having at least one closely matching network resource that may be used as a redundant network resource. (Specification p. 5, ¶ [13], ll. 2-4). The method comprises receiving a request for an execution of a task on one of said plurality of network resources. (Specification p.5, ¶ [13], ll. 4-5 and p. 12, ¶ [33], ll. 2-3 and Fig. 2A item S210). The least loaded of at least one network resource from said plurality of network resources or at least one matching redundant resource from said plurality of network resources is assigned to the request for execution of the task. (Specification p.5, ¶ [13], ll. 6-7 and p. 13, ¶ [34], ll. 1-4 and Fig. 2B, item S250). One of said plurality of communication paths or one of the plurality of redundant communication paths is assigned to the request for execution of the task. (Specification p. 5, ¶ [13], ll. 7-8 and p. 12, ¶ [33], ll.5-9 and Fig. 2A, items S215 and S220). The requestor of the

assigned network resource and the assigned communication path is informed. (Specification p.5, ¶ [13], ll. 8-9 and p. 12, ¶ [33], ll.5-9 and Fig. 2A, S225).

Independent claim 21 recites a method for re-balancing execution loads in a network system. (Specification p. 5, ¶ [14], ll. 1-2). The network contains a plurality of communication paths, a plurality of redundant communication paths independent of said plurality of communication paths, a plurality of network resources of differing types, each of said network resources having at least one closely matching network resource which may be used as a redundant network resource. (Specification p. 5, ¶ [14], ll. 2-4 and p. 14, ¶ [36], ll.5-9 and Fig. 3A, item S310). The method comprises determining the type of failure that caused the failure notification. (Specification p. 5, ¶ [14], ll. 4-6, and Fig. 3A, item S315). If a communication path has failed, and if no alternative communication path is available, an error notification is issued. (Specification p. 5, ¶ [14], ll. 6-7 and p. 14, ¶ [36], ll. 4-10 and Fig. 3A, items S315, S320, S325). If an alternative communication path is available, the failed communication path is eliminated from further use, and the load is redistributed. (Specification p. 5, ¶ [14], ll. 7-9 and p. 15, ¶ [39], ll. 2-5 and Fig. 3C, items S355, S360). If a network resource has failed, and no alternative network resource is available, an error notification is issued. (Specification p. 5, ¶ [14], ll. 9-11 and p. 14, ¶ [37], ll.1-6 and Fig. 3B, items S315 and S335). If an alternative network resource is available, the failed network resource is eliminated from further use, and the load is redistributed. (Specification p. 5, ¶ [14], ll. 11-13 and p. 14, ¶ [38], ll.1-5 and Fig. 3B, items S345 and S350).

Independent claim 22 recites a computer software product for balancing loads in a network system. (Specification p. 6, ¶ [15], ll. 1-2 and p. 16, ¶ [41], ll. 1-3). The network system contains a plurality of communication paths, a plurality of redundant communication paths, a plurality of network resources of differing types and a plurality of redundant network resources. (Specification p. 6, ¶ [15], ll. 2-5 and p. 16, ¶ [41], ll. 3-5). The computer program product comprises software instructions for enabling the network system to perform predetermined operations, and a computer readable medium bearing the software instructions. (Specification p. 6, ¶ [15], ll. 5-7 and p. 16, ¶ [41], ll. 5-8). The predetermined operations comprise receiving a request for access to one of the plurality of network resources, and assigning at least one network resource ore at least one matching redundant resource from the plurality of network resources to the request for executing the task. (Specification p. 6, ¶ [15], ll. 7-10 and p. 16, ¶ [41], ll. 8-14). The predetermined operations further comprise assigning one of said plurality of communication paths or one of said plurality of redundant communication paths to the request for executing the task, and informing the requestor of the assigned network resource and the assigned communication path. (Specification p. 6, ¶ [15], ll. 10-13 and p. 16, ¶ [41], ll. 14-17).

Independent claim 27 recites a computer software product for re-balancing loads in a network system. (Specification p. 6, ¶ [16], ll. 1-2 and p. 16, ¶ [42], ll. 1-3). The network contains a plurality of communication paths, a plurality of redundant communication paths, a plurality of network resources of differing types and a plurality of redundant network resources. (Specification p. 6, ¶ [16], ll. 2-5 and p. 16, ¶ [42], ll. 2-6). The computer program product

comprises software instructions for enabling the network system to perform predetermined operations and a computer readable medium bearing the software instructions. (Specification p. 6, ¶ [16], ll. 5-7 and p. 16, ¶ [42], ll. 5-8). The predetermined operations comprise determining the type of failure that caused the failure notification. (Specification p. 6, ¶ [16], ll. 7-9 and p. 16, ¶ [42], ll. 8-9). If a communication path has failed, and if no alternative communication path is available, the predetermined operations issue an error notification. (Specification p. 6, ¶ [16], ll. 9-11 and p. 16, ¶ [42], ll. 9-12). If an alternative communication path is available, the predetermined operations eliminate the failed communication path from further use, and the load is redistributed. (Specification p. 6, ¶ [16], ll. 11-13 and p. 16, ¶ [42], ll. 12-15). If a network resource has failed, and if no alternative network resource is available, the predetermined operations issue an error notification. (Specification p. 6, ¶ [16], ll. 13-14 and p. 16, ¶ [42], ll. 15-18). If an alternative network resource is available, the predetermined operations eliminate the failed network resource from further use, and the load is redistributed. (Specification p. 6, ¶ [16], ll. 14-17 and p. 16, ¶ [41], ll. 18-21).

Independent claim 28 recites a redundant network system capable of using redundant elements for the purpose of load balancing. (Specification p. 7, ¶ [17], ll. 1-3). The system comprises at least one client node and at least two network switches providing alternate connection paths to the client node. (Specification p. 7, ¶ [17], ll. 3-4 and p. 20, ¶ [47], ll. 4-6). The system further comprises at least two cache control nodes capable of supporting an address resolution protocol and capable of load balancing storage control nodes. (Specification p. 7, ¶ [17], ll. 4-7 and p. 20, ¶ [47], ll. 7-9). The cache control nodes connected to the network

switches. (Specification p. 7, ¶ [17], ll. 7-8 and p. 20, ¶ [47], ll. 8-11). The system further comprises at least two storage control nodes and the storage control nodes connected to at least the network switches. (Specification p. 7, ¶ [17], ll. 8-10 and p. 20, ¶ [47], ll. 6-8).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The issues on appeal are whether the following rejections are proper:

1. Rejection of claims 1-3, 5-7, 9, 10, 16 and 22 under 35 U.S.C. 103(a) as being unpatentable over Wolff in view of Peterson.
2. Rejection of claims 4, 8, 11, 12, 13, 14, 15, 17-20, 23-26 and 28-38 under 35 U.S.C. 103(a) as being unpatentable over Wolff in view of Peterson et al. in further view of Richter.
3. Rejection of claims 21 and 27 under 35 U.S.C. 103(a) as being unpatentable over Wolff in view of Richter.

VII. ARGUMENT

1. Claims 1-3, 5-7, 9, 10, 16 and 22 are unobvious over the combined teachings of Wolff and Peterson.

The present invention, as recited in claim 1, requires at least one redundant matching resource. Further, a computer is required that transfers tasks from execution on the network resource for execution to the redundant matching resource if the network resource fails. The computer balances execution loads between the tasks executed by the network resource and the tasks executed by the redundant matching resource. Further, a communication medium is required that connects the computer, the terminal node, the network resource and the redundant matching resource. The communication medium is required to have at least two independent communication paths between the terminal node, the network resource and the redundant matching resource.

The present invention overcomes a deficiency in the prior art where a redundant system would idle until it is utilized due to a failure. The present invention provides the capability to have different tasks operating on both matching network resources, balancing the load between the two. In addition, in case of a failure, the tasks are transferred from the failing networked resource to the other resource.

The Appellants respectfully submit that, Wolff fails to teach or suggest redundancy of systems. The patents office's finding that Wolff teaches this feature is believed to be at least due to the Patent Office's erroneous interpretation of load-balancing activity and redundancy. The computer in claim 1 is required to performs load-balancing between tasks given to a network

resource and tasks given to its redundant network resource for the purpose of utilizing also the redundant resources. These tasks are not necessarily the same tasks. The computer is also required to move tasks from a failing network resource to the matching redundant resource. Wolff does not suggest this capability.

The present invention, as recited in claim 1 also requires the communication medium to have at least two independent communication paths between the terminal node, the network resource and the redundant matching resource. Wolff arguably discloses the remapping of network paths. However this is believed to be completely different from having independent pathways of performing the communication. For example, the case of a network where it is possible to get from A to E in the following ways: $A \rightarrow B \rightarrow C \rightarrow E$ or alternatively, remapping in accordance with Wolff, by means of pathway $A \rightarrow D \rightarrow C \rightarrow E$. Such a remapping may enable an alternate path, but such a path is not redundant, because if the $C \rightarrow E$ network fails, the choice of the alternate path will not help. In the system suggested by Wolff, no redundant pathway exists between A and E.

The above feature of Wolff merely provides the ability to remap a communication path. However, such an ability to remap a communication path cannot be construed to be a suggestion to map a **redundant** communication path. To map such a redundant route as in the present invention, will require at least a redundant network path that is completely independent from the original network path. Wolff does not suggest such an independent path. Regardless of how intelligent the re-routing system of Wolff is, without a completely independent communication path the features of the present invention cannot be achieved.

In the Office Action dated June 25, 2005, the Patent Office acknowledges that Wolff fails to disclose, *inter alia*, at least one matching resource. The Patent Office alleges that Peterson *et al.* provides the necessary disclosure to overcome the acknowledged deficiencies of Wolff. The Patent Office contends that Peterson *et al.* disclose a mirror system which the Patent Office erroneously construes to be a redundant system as in the present invention. Such a mirror system could be considered as data redundant, a skilled artisan noticing that the difference between the invented system and Peterson *et al.* is that in the invented system different tasks are executed on each of the systems, while in the mirrored system, at all times, the data on the main and mirrored resources is required to be identical. This is at least because the mirror systems are by definition to be duplicates of one another. Therefore, the mirror systems are, by definition, data balanced. Typically, they will not require further balancing.

In fact, such a mirrored memory is clearly defined by Peterson to be something completely different from the present invention. Peterson notes that "...mirrored memory simply means that data in one memory is duplicated or 'mirrored' in another memory." (column 2, line 65 to column 3, line 1). In other words, data that is on one system is copied exactly onto the other system. Peterson does not provide any suggestion that a mirrored system can provide the capabilities of load-balancing between two identical systems. In fact, providing such a functionality is believed to interfere with the well-known functions of the mirrored systems.

The combination of Wolff and Peterson *et al.* fails to teach or suggest at least a computer that balances task loads between resources and matching redundant resources using both a communication medium and a redundant communication medium, as recited in amended claim

1. As discussed above, the combination of references does not teach or suggest at least load-balancing between tasks given to a network resource and tasks given to its redundant network resource, nor does the combination of references teach or suggest at least a redundant communication medium that is separate from a communications medium coupled to a network resource, as recited in claim 1.

Wolff provides the ability of balancing loads between a plurality of networked resources. However, Wolff does not suggest redundant portions of the system that are normally kept idle awaiting a failure in the system. The fact that an alternate path exists in a networked system does not suggest that there is a redundant path. The addition of a mirror system as in Peterson merely achieves that both systems have the identical data in any given period of time. Therefore, the combined teaching of Wolff in view of Peterson does not amount to the invention disclosed in the Application.

To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest **all the claim limitations**. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. MPEP 2142 citing *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). See MPEP § 2143 - § 2143.03 for decisions pertinent to each of these criteria.

Clearly, the Examiner has not satisfied the “all limitations” prong to establish a *prima facie* case of obviousness based on the combined teachings of Wolff and Peterson *et al.*

Since neither Wolff nor Peterson *et al.* teach or suggest at least the load-balancing and redundant communications medium recited in amended claim 1, Applicant submit that one of skill in the art would not be motivated to combine the references.

Further, the Patent Office finds alleged motivation to combine references in the need to provide a fault tolerant environment. However, the present invention is believed to go far beyond providing simple fault tolerant systems and provides techniques for using a network resource and its redundant network resource in a load-balanced fashion rather than leaving the redundant resource idling until such time that it is required. Thus, Appellants submit that the combination of Wolff and Peterson *et al.* fail to meet the motivation prong of a *prima facie* case of obviousness with respect to claim 1.

Therefore, the Appellants respectfully submit that, the patent office’s finding of obviousness of claim 1 based on Wolff and Peterson must be reversed.

Claims 9, 16 and 22 includes limitations analogous to the ones discussed above in relation to claim 1. Therefore, the Appellants respectfully submit that, the patent office’s finding of obviousness of claims 9, 16 and 22 based on Wolff and Peterson must be reversed.

Claims 2, 3, 5-7 are allowable as well, at least by virtue of their dependency from claim 1. Appellants respectfully request that the Patent Office’s § 103(a) rejection of claims 1-7 be reversed.

Claim 10 is allowable as well, at least by virtue of its dependency from claim 9.

Appellants respectfully request that the Patent Office's § 103(a) rejection of claim 20 be reversed.

2. Claims 21 and 27 are unobvious over the combined teachings of Wolff and Richter.

The combined teachings of Wolff and Richter *et al.* do not suggest at least load redistribution between remaining communication paths and remaining redundant communication paths, as recited in independent claims 21 and 27. As discussed above with respect to the other independent claims, neither Wolff nor Richter *et al.* teach or suggest at least the load balancing and redundant communication path assignment features of the present invention recited in claims 21 and 27. Thus, Appellants submit that the combination of Wolff and Richter *et al.* fail to meet the “all limitations” prong of a *prima facie* case of obviousness with respect to claims 21 and 27.

At least because of the above-noted deficiencies in the combined teachings of Wolff and Richter *et al.* the Appellants submit that one of skill in the art would not be motivated to combine the references. Thus, Appellants submit that the combination of Wolff and Richter *et al.* fail to meet the motivation prong of a *prima facie* case of obviousness with respect to claims 21 and 27. Therefore, the patent office’s finding of obviousness of claims 21 and 27 based on the combined teachings of Wolff and Richter must be reversed.

3. Claims 4, 8, 11, 12, 13, 14, 15, 17-20, 23-26 and 28-38 are unobvious over the combined teachings of Wolff, Peterson and Richter.

Claim 28 includes limitations analogous to claim 1 and should be allowed at least for analogous reasons.

Claims 4, 8, 11, 12, 13, 14, 15, 17-20, 23-26 and 29-38 are dependent on base claims 1, 9, 16, 22, 27 and 28 and are allowable for at least the same reasons.

Further, Richter does not overcome the deficiency noted above in the combined teachings of Wolff and Peterson.

The combined teachings on Wolff, Peterson and Richter continue to fail to meet the “all limitations” prong required to establish *prima facie* obviousness. Therefore, the patent office’s finding of obviousness of the above claims based on the combined teachings of Wolff, Peterson and Richter must be reversed.

VIII. CONCLUSORY REMARKS

Unless a check is submitted herewith for the fee required under 37 C.F.R. §41.37(a) and 1.17(c), please charge said fee to Deposit Account No. 19-4880.

The USPTO is directed and authorized to charge all required fees, except for the Issue Fee and the Publication Fee, to Deposit Account No. 19-4880. Please also credit any overpayments to said Deposit Account.

Respectfully submitted,

SUGHRUE MION, PLLC
Telephone: (202) 293-7060
Facsimile: (202) 293-7860

WASHINGTON OFFICE

23373

CUSTOMER NUMBER

Chid S. Iyer
Registration No. 43,355

Date: October 16, 2006

CLAIMS APPENDIX

CLAIMS 1-38 ON APPEAL:

1. (previously presented): A system comprising:

at least one terminal node;

at least one network resource, said network resource having at least one redundant matching resource;

a computer that transfers tasks from execution on said network resource for execution on said redundant matching resource if said network resource fails, and that balances execution loads between the tasks executed by said network resource and the tasks executed by said redundant matching resource; and

a communication medium connecting said computer, said terminal node, said network resource and said redundant matching resource, said communication medium having at least two independent communication paths between said terminal node, said network resource and said redundant matching resource.

2. (original): The system of claim 1, wherein said communication medium is selected from a group consisting of a local area network (LAN), a wide area network (WAN), an Ethernet based network, an Internet protocol (IP) based network, an asynchronous transfer mode (ATM)

based network, a peripheral component interconnect (PCI) based network, and an InfiniBand based network.

3. (original): The system of claim 1, wherein said communication medium further comprises at least one communication element.

4. (original): The system of claim 3, wherein said at least one communication element is selected from a group consisting of a network switch and a cache control node.

5. (previously presented): The system of claim 1, wherein said at least one network resource is selected from a group consisting of a storage device, a redundant array of independent disks, a cache, a file system, and a location independent file system.

6. (previously presented): The system of claim 1, wherein said at least one redundant matching resource is selected from a group consisting of a storage device, a redundant array of independent disks, a file system, and a location independent file system.

7. (original): The system of claim 1, wherein said at least one network resource and said at least one redundant matching resource are continuously interchangeable.

8. (original): The system of claim 1, wherein said computer load balances by concurrently assigning tasks to said at least one network resource and said at least one redundant matching resource.

9. (previously presented): A system comprising:

- a plurality of terminal nodes;
- a plurality of network resources;
- a plurality of redundant resources, each of said plurality of network resources closely matching at least one of said plurality of redundant resources;
- a computer that moves tasks from a failed network resource to a redundant resource that closely matches the failed network resources, and that balances execution loads between the tasks executed by said plurality of network resources and the tasks executed by said plurality of redundant resources; and
- a communication medium connecting said computer, said terminal nodes, said network resources and said redundant resources, said communication medium having at least two independent communication paths between said terminal nodes, said network resources and said redundant resources.

10. (previously presented): The system of claim 9, wherein said computer, after receiving a request for using said network resources, balances the load by:

assigning the least loaded of at least one network resource from said plurality of network resources or at least one matching redundant resource from said plurality of network resources to the request;

determining the least loaded communication path in said communication medium; and

informing the requestor of the assigned network resource and the determined communication path.

11. (previously presented): The system of claim 10, wherein the assigning of at least one network resource comprises:

if the number of available network resources is larger than one, executing a selection function to determine the at least one network resource to be assigned, and assigning that resource to the requestor; and

if the number of available network resources is equal to one, assigning the available resource to the requestor.

12. (original): The system of claim 11, wherein the selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

13. (previously presented): The system of claim 10, wherein determining a communication path in said communication medium comprises:

if the number of available communication paths is larger than one, executing a selection function to determine a communication path to be used, and informing requestor of selected communication path; and

if the number of available communication paths is equal to one, informing requestor of available communication path.

14.: (original): The system of claim 13, wherein the selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

15. (original): The system of claim 9, wherein, if said computer receives a failure notification, said computer attempts to re-balance loads by:

determining the type of failure that caused the failure notification;

if the computer determines that a communication path in said communication medium has failed, then:

if no alternative communication path is available, the computer issues an error notification;

otherwise, the failed communication path is eliminated from further use, and the load is redistributed;

if the computer determines that a network resource has failed, then:

if no alternative network resource is available, the computer issues an error notification;

otherwise, the failed network resource is eliminated from further use, and the load is redistributed.

16. (previously presented): A method for balancing loads of tasks executed in a network system containing a plurality of communication paths, a plurality of redundant communication paths independent of said plurality of communication paths, a plurality of network resources of differing types each of said network resources having at least one closely matching network resource that may be used as a redundant network resource, wherein the method comprises:

receiving a request for an execution of a task on one of said plurality of network resources;

assigning, for a purpose of an execution of the task, the least loaded of at least one network resource from said plurality of network resources or at least one matching redundant resource from said plurality of network resources to the request;

assigning, for a purpose of an execution of the task, one of said plurality of communication paths or one of the plurality of redundant communication paths to the request; and

informing the requestor of the assigned network resource and the assigned communication path.

17. (original): The method of claim 16, wherein the assigning of at least one network resource comprises:

determining the least loaded network resource of the requested resource type;

if the number of available network resources is larger than one, executing a selection function to determine the at least one network resource to be assigned, and assigning that network resource to the requestor; and

if the number of available network resources is equal to one, assigning the available network resource to the requestor.

18. (original): The method of claim 17, wherein said selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

19. (original): The method of claim 16, wherein assigning a communication path comprises:

determining the least loaded communication path to said assigned network resource;

if the number of available communication paths is larger than one, executing a selection function to determine a communication path to be used, and informing requestor of selected communication path; and

if the number of available communication paths is equal to one, informing requestor of available communication path.

20. (original): The method of claim 19, wherein said selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

21. (previously presented): A method for re-balancing execution loads in a network system containing a plurality of communication paths, a plurality of redundant communication

paths independent of said plurality of communication paths, a plurality of network resources of differing types, each of said network resources having at least one closely matching network resource which may be used as a redundant network resource, wherein said method comprises:

determining the type of failure that caused the failure notification;

if a communication path has failed, then:

if no alternative communication path is available, issuing an error notification;

otherwise, the failed communication path is eliminated from further use, and the load is redistributed between the remaining communication paths and the remaining independent redundant communication paths;

if a network resource has failed, then:

if no alternative network resource is available, issuing an error notification;

otherwise, the failed network resource is eliminated from further use, and the load is redistributed between the remaining network resources and the remaining redundant network resources.

22. (previously presented): A computer software product for balancing execution loads in a network system containing a plurality of communication paths, a plurality of redundant communication paths independent of said plurality of communication paths, a plurality of network resources of differing types, each of said network resources having at least one closely

matching network resource which may be used as a redundant network resource, the computer program product comprising:

software instructions for enabling the network system to perform predetermined operations, and a computer readable medium bearing the software instructions, said predetermined operations comprising:

receiving a request for access to one of said plurality of network resources;

assigning, for a purpose of an execution of the task, the least loaded of at least one network resource from said plurality of network resources or at least one matching redundant resource from said plurality of network resources to the request;

assigning, for a purpose of an execution of the task, one of said plurality of communication paths to the request or one of said plurality of redundant communication paths;
and

informing the requestor of the assigned network resource and the assigned communication path.

23. (original): The computer software product of claim 22, wherein assigning of at least one network resource comprises:

determining the least loaded network resource of the requested resource type;

if the number of available network resources is larger than one, executing a selection function to determine the at least one network resource to be assigned, and assigning that network resource to the requestor; and

if the number of available network resources is equal to one, assigning the available network resource to the requestor.

24. (original): The computer software product of claim 23, wherein said selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

25. (original): The computer software product of claim 22, wherein assigning a communication path comprises:

determining the least loaded communication path to said assigned network resource;

if the number of available communication paths is larger than one, executing a selection function to determine a communication path to be used, and informing requestor of selected communication path; and

if the number of available communication paths is equal to one, informing requestor of available communication path.

26. (original): The computer software product of claim 25, wherein said selection function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

27. (previously presented): A computer software product for re-balancing execution loads in a network system containing a plurality of communication paths, a plurality of redundant communication paths independent of said plurality of said communication paths, a plurality of network resources of differing types, each of said network resources having at least one closely matching network resource which may be used as a redundant network resource, the computer program product comprising:

software instructions for enabling the network system to perform predetermined operations, and a computer readable medium bearing the software instructions, the predetermined operations comprising:

determining the type of failure that caused the failure notification;

if a communication path has failed, then:

if no alternative communication path is available, issuing an error notification;

otherwise, the failed communication path is eliminated from further use, and the load is redistributed between the remaining communication paths and the remaining independent redundant communication paths;

if a network resource has failed, then:

if no alternative network resource is available, issuing an error notification;

otherwise, the failed network resource is eliminated from further use, and the load is redistributed between the remaining network resources and the remaining redundant network resources.

28. (previously presented): A redundant network system capable of using redundant elements for the purpose of load balancing the system comprising:

at least one client node;

at least two network switches providing alternate connection paths to said client node, said connection paths being independent from each other;

at least two cache control nodes capable of supporting an address resolution protocol and capable of load balancing storage control nodes, said cache control nodes connected to said network switches; and

at least two storage control nodes, said storage control nodes connected to at least said network switches.

29. (original): The system of claim 28, wherein said address resolution protocol is executed on one of said cache control nodes.

30. (original): The system of claim 28, wherein one of said cache control nodes is used as a redundant cache control node.

31. (original): The system of claim 30, wherein the second cache control node receives address resolution protocol information from the first cache control node.

32. (original): The system of claim 31, wherein, when one of said cache control nodes fails, the remaining cache control node executes said address resolution protocol.

33. (original): The system of claim 28, wherein one of said cache control nodes generates a media access control address corresponding to a storage control node.

34. (original): The system of claim 33, wherein said media access control address is for the storage control node which is the next to receive the access request based on a load balance function executed by one of the cache control nodes.

35. (original): The system of claim 34, wherein said load balancing function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

36. (original): The system of claim 34, wherein said media access control address is further based on the specific network path to be used for said client node to connect to said storage control node.

37. (original): The system of claim 33, wherein said media access control address is for the network path which is the next to receive the access request based on a load balance function executed by one of the cache control nodes.

38. (original): The system of claim 37, wherein said load balancing function is selected from a group comprising a round robin function, a weighted round robin function, a random function, a least loaded function and a least recently used function.

EVIDENCE APPENDIX:

None.

RELATED PROCEEDINGS APPENDIX

None.